Bolt Beranek and Newman Inc.

LEVEL II

Report No. 3961

Command and Control Related Computer Technology: Packet Radio

Quarterly Progress Report No. 15
1 June 1978 to 31 August 1978

January 1979

Prepared for:
Defense Advanced Research Projects Agency

(14) BBN [redacted]-3961

(11) Jan[redacted]79

(12) 43 p.

(6) COMMAND AND CONTROL RELATED COMPUTER TECHNOLOGY:
Packet Radio.

(9) Quarterly Progress Report. No. 15,
1 Jun[redacted]31 Aug[redacted]78,

(10)
M. /Beeler,
J. /Burchfiel,
R. /Perlman,
W. W. /Plummer
V. /Strazisar

This research was supported by the Defense Advanced Research Projects Agency under ARPA Order No. 2935 Contract No. MDA903-75-C-0180,

(15) ARPA Order-2935

060 100

79 01 23 493

mt

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 3961	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) COMMAND AND CONTROL RELATED COMPUTER TECHNOLOGY: PACKET RADIO		5. TYPE OF REPORT & PERIOD COVERED Quarterly Progress Report 1Jun.78-31Aug.78
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) M.Beeler W.W.Plummer J.Hirschkorn J.Burchfiel V.Strazisar R.Perlman J.Sussman		8. CONTRACT OR GRANT NUMBER(s) MDA903-75-C-0180
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman, Inc. 50 Moulton Street Cambridge, Massachusetts 02138		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency Information Processing Techniques Office 1400 Wilson Blvd, Arlington VA 22209		12. REPORT DATE January 1979
		13. NUMBER OF PAGES 39
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Distribution of this document is unlimited. It may be released to the Clearinghouse; Department of Commerce for sale to the general public.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

This research was supported by the Defense Advanced Research Projects Agency under ARPA Order No. 2935.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

packet radio, computer communications, PDP-11, TCP, station gateway

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This document describes progress on development of a packet radio network. Activities reported include work on Station Software and Internetworking Research and Development.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

TABLE OF CONTENTS

1. INTRODUCTION

The Packet Radio project relies heavily on station software for a variety of control, coordination and monitoring functions. The role of BBN in developing this software is to specify, design, implement and deliver programs which implement these functions.

During this quarter significant progress on point-to-point routing and Packet Radio Net (PRN) performance was made, as well as important advances in PRN design and in Internet facilities. Section 2 of this report covers progress in design, documentation, and resolutions reached jointly with other Packet Radio contractors. The major design work focused on two areas identified by ARPA at the previous PR Working Group meeting as high priority. These are routing compatible with a stationless area of an operating net, and specification of a local connectivity monitoring mechanism (LROPs) and associated reporting of the results to the station.

Section 3 covers the delivery of station software to fully use the point-to-point routing capability of the PR units. Additionally, testing activities performed this quarter are discussed in Section 3. These tests are both to verify correct operation of the point-to-point software and to gather data on performance of the PR net as a whole. Reports from UCLA of lost cumulative statistics (CUMSTAT) packets during measurement experiment runs were a special target in this investigation. Efforts in support of station software are also reported in Section 3.

Section 4 deals with progress in the Internet arena. There are two important aspects to the Internet work performed this quarter. First, version 2.5 of the Transmission Control Program was released to the SRI-KA site. Second, a vigorous attack on

the performance of the Very Distant Host (VDH) interface, employed in this case between gateways and Satellite network IMPs (SIMPs), was mounted. These efforts have led to the design, implementation and testing of modifications to the VDH code to improve performance.

Section 5 deals with hardware issues. During this quarter, the major progress in the hardware domain has been site preparations for the Boston area PR network. A number of troublesome details came to light only as work progressed, but no major pitfalls were encountered.

2. MEETINGS, TRIPS, PUBLICATIONS

2.1. Meetings and Trips

BBN hosted a meeting of the Packet Radio Working Group during June 19-21 this quarter. We presented our status and plans, and discussed potential problem areas as usual; but with representatives in personal attendance at our facilities, we were additionally able to demonstrate the new point-to-point labeler. This was especially valuable as a familiarization activity for SRI personnel, allowing them to see the software delivered later in the quarter with our staff in attendance to answer questions and illustrate features. Other areas in which fruitful negotiations of especial relevance to efforts at BBN were held at the meeting are 1822 interface design, and the need for user documentation of PR net protocols.

At the end of June, Julie Sussman of our staff visited SRI for the purpose of testing the new station software with point-to-point routing, particularly the labeler module. The information gained on this visit did aid in final debugging of the labeler, but the principal return was not expected; the field PR net proved insufficiently reliable to obtain solid experimental results on the performance of the new labeler. This points toward a need for more extensive testing and debugging of network components, particularly the hardware, but not to the exclusion of concern over software testing as well. A detailed discussion of the results of this trip appears in Section 3.

Virginia Strazisar and Radia Perlman attended the Internet and Packet Satellite Working Group meetings held at Lincoln Labs the first week of August. Virginia led a small working group in discussing error reporting in the catenet. The types of errors that the group thought should be reported were: host dead or unreachable, network unreachable, and type of service

unavailable. There was some discussion on exactly what information should be included in an error report. Although no definite format was defined for error reports, the group felt that there must be sufficient information for the source of the packet to identify the segment that caused the error. The details of an internet error reporting facility must still be specified and implemented.

During this quarter BBN personnel attended the TCP meeting at MIT, June 15-16. BBN personnel were also on hand at SRI during the week of June 5 for the BAA demonstration. Although little arose requiring our attention there, our on-site attendance extended through the demonstration.

2.2. Publications

PRTN 174 - revision 6, "Packet Radio Network Station Labeling Process"

PRTN 212 - revision 5, "Specification of Measurement File Entries"

Revised sections of **station operator's guide**: table of contents, station overview, and chapters on the connection, labeler, gateway and measurement processes.

These two updated PRTNs, plus the guide chapters constitute the documentation package for the point-to-point routing station software released this quarter. The PRTNs were sent by U.S. mail, as usual; the operator's guide chapters were placed on-line on [BBNA], as usual.

PRTN 255, "LROPs and Neighbor Tables"

As discussed in last quarter's report, this PRTN provides a detailed technical discussion of design issues and alternative solutions for the new mechanism to collect network connectivity data. The actual mailing of this PRTN occurred early in this quarter, providing a detailed focus for LROP discussion at the

PRWG meeting at BBN later in June, led by the author of PRTN 255.

PRTN 256, "Stationless Compatible PR NET Routing"

In PRTN 256 a number of complexly related issues are treated. The main problem is design of routing for a PR net in which some areas are under control of a station, while PRs in other areas are not. Thus, the latter areas are stationless, at least temporarily; and the routing algorithm must operate not only in each kind of area but also across the boundary. The most difficult part of the solution design is ensuring the compatibility between stationless and stationful areas, so new routes crossing such boundaries can be established and used. This PRTN examines in detail the design we propose. Subtleties arise in making sure such secondary aspects of routing as alternate routing, hop acknowledgments and route failure notification will work. We continue our long-standing concern with efficiency and robustness of PRN protocols in sections on loop detection and detection of excessive alternate routing. Contents of packet headers, routing control packets and routing tables in PRs are listed; algorithms for the PR to use are explained concisely but in detail.

Internet Experimenters Note (IEN) 45, "TCP Checksum Function Design"

This paper discusses various desiderata in choosing a function to provide error detection in a Transmission Control Program, and evaluates various candidate functions in the light of these criteria. The paper thus helps remove the mystery from this aspect of TCP design and implementation, by documenting the considerations usually pursued only informally. Prior to its release as an IEN, this paper was distributed at the TCP meeting at MIT in June.

2.3. Negotiations and Informal Documents

2.3.1. Changes in measurement file

Development of the point-to-point routing labeler (section 3.1.2) included changes to entries made by the labeler in the measurement file. The cumulative statistics entry was augmented to include counts related to point-to-point route assignment, and a new entry was created to report point-to-point routes. In keeping with our usual policy, we sent UCLA an advance copy of the changes to PRTN 212, defining these entries, so that they could get a head start on modifying their programs to match.

2.3.2. LROPs and neighbor tables

We sent PRSETD a summary of ways in which the design of LROPs and neighbor tables went beyond PRTN 255 in discussions at the June 19-21 Packet Radio Working Group meeting. In particular, an explicit division (incremental packet receive count divided by incremental packet transmit count, R/T) to obtain link quality is not necessary; alternative algebraic approaches are available. This is good because performing a true division in the EPR's CPU, an IMP-16L, is slow and difficult. On the subject of station request PDPs, Collins stated that a station will not have to have labeled a PR for the PR to honor the station's request for a Performance Data Packet. A main topic which received considerable attention is the problem of rapidly servicing the changing connectivity of a mobile PR. One possibility is for the station to have a way to tell a PR not to overwrite a given neighbor table slot. This could be used to "lock in" slots relating to a mobile PR until the PR had left the vicinity. The alternative we favor is that proposed in PRTN 255, the Distress PDP, which roughly corresponds to today's Distress ROP (which will be eliminated in the LROP/neighbor table scheme). In both the Distress ROP and the Distress PDP, the idea is to use

a message, carried over the first hop by broadcast protocol, to immediately inform the station that no hop acknowledgement could be obtained for a packet. Whether this Distress PDP (and suitable follow-up action taken by the station) is adequate to quickly reestablish traffic flow to a mobile PR was not obvious at the meeting. Thus, this issue deserves further study.

2.3.3. Symmetrical 1822

Also prepared as a result of discussions at the PRWG meeting was a detailed list of issues resolved regarding a "symmetrical" 1822 interface. The need for specification of such an interface between PRs and attached devices has prompted considerable design and negotiation efforts at BBN, peaking with publication of PRTN 245 last quarter. This has brought the issues into clear focus, and a working session at the June 19-21 PRWG meeting arrived at resolutions on all eleven outstanding issues. In most cases the resolution is complete and final; in a few, specific steps or testing was agreed upon to reach a final agreement.

2.3.4. Stationless compatible routing

Also as an outcome of the PRWG meeting, we summarized discussion of PRTN 256, on stationless compatible routing. The main issue which emerged is the choice of storing routing table entries indexed by destination only, or by source/destination pair. It was resolved that BBN will prepare a PRTN comparing and contrasting these two.

2.3.5. Residual CAP4 issues

A final informal document to arise from the meeting was a summary of residual issues in the CAP4 protocol. Two issues bearing on longer term efforts were the report by SRI of serious degradation of apparent link quality when one of a PR's six buffers is tied up, and report by UCLA of improper delivery of

cumulative statistics (CUMSTAT) packets during measurement experiment runs. The former of these is leading to drafting a more explicit note of concern at BBN as this quarter ends; the latter has led to extensive testing and analysis of network operation during this quarter, as detailed in section 3.

Several decisions about CAP4 protocol changes were reached at the PRWG meeting this quarter, and further discussion of unresolved issues took place after the meeting. These CAP4 items are as follows:

A. The two changes we had proposed last quarter to DROP (Distress ROP) handling will be implemented. These changes, to enable DROPs in all PRs (not just terminals) and to avoid broadcasting of DROPs in favor of direct routing when possible, were described in detail in our last quarterly report. We pointed out that if this DROP delivery change was implemented without a slight change in the station labeler program, that the labeler might report that the PR emitting the DROP had connectivity to itself. This would in no way affect station operation, but could be confusing to a station operator.

B. In last quarter's report we described in detail some complex protocol issues concerning duplicate detection, recognition of hop acknowledgements, alternate routing, and folded routes (routes which pass through a PR twice). We pointed out that the protocol being implemented was not bug-free, but would have problems in two very unlikely circumstances, namely with certain folded routes and with copies of an alternate-routed packet which take different length paths. It was agreed that the protocol change discussed in the report, to make use of hop counts in duplicate filtering, would be implemented.

C. SRI felt that it would be desirable to have a way of telling a PR to "shut up and listen". One problem is that if the station totally shuts up a PR, then if the station is reinitialized it will not know about the PR because it will not hear ROPs from it. It was decided that initially turning PRs off and on will be done from TIUs with a command packet. Station involvement will be reconsidered later.

D. SRI wanted a means of setting up particular point-to-point routes for experiments. One suggestion was an "insert route" command packet that would be used from a TIU. The preferred suggestion was that the station be able to maintain whatever routes the operator desires rather than ones it would choose itself. We agreed to look into the feasibility of implementing this, but station space constraints make it unlikely that we could do so immediately.

E. Some net test situations (such as SRI's PMON program) call for forwarding by the station rather than use of point-to-point routes. A TIU can force override of the PTP routes, but the station, seeing the packets to forward, will keep assigning PTP routes. Two ideas to prevent this were discussed: a reflector in the station, to return a packet to its source; a new header bit saying "don't give me a PTP route". This is discussed more fully below.

F. Routes in packets in the net will be in a standard form, always being read in the same direction. (Currently a route can be read either way depending on the direction bit.) A PR wishing to reverse a packet can do so by reversing the actual route. An end device wishing to reverse a packet can clear the direction bit; the attached PR will reverse the route and set the bit. This removes the current asymmetry in alternate-routing which could expand a route only if it was

being processed forwards. It was originally thought that this change would have no effect on the station or TIUs, but we subsequently realized there would be an adverse effect on point-to-point route handling.

Until now, whenever a PR inserted a route in a packet, the PR itself appeared first and the route was incrementing. Decrement packets were proceeding from a PR that had not inserted the route toward the first PR in the route, namely the one where the route had been inserted. With this change, however, all packets in the net, whether or not they are coming from the PR which inserted the route, are incrementing, and the first PR in the route needn't be the one that inserted the route.

For PTP route assignment it is of interest to know who inserted the route. If the station receives a packet which it has to forward, then it would rather not forward it but instead assign a PTP route. The only place it is useful to try to store such a route is in the PR which is inserting the route, which is currently sending the packet to the station for forwarding. In the current CAP, if the station sees a decrementing packet, it knows it did not get there from a PR that put a route in it, so there is no reason to try to assign a PTP route. For an incrementing packet, it knows the first PR in the route would benefit from having a PTP route, and tries to assign one.

With the proposed change the station can't distinguish these cases. Thus if A is talking to D and inserts the route A-B-S(tation), then the packet is forwarded S-C-D and a PTP route will be given to A to get to D. Then if D turns the packet around (e.g. to ACK it), D-C-S heading for A, the station will forward it and try to give D a PTP route to get

to A. Since D didn't do any route lookup, it didn't reserve a table slot, so it will not store the route (which would not be used anyway).

Now, if a route succeeds in getting stored in A (say A-X-Y-D) then future packets reversed by D will be D-Y-X-A and will not involve the station or cause any trouble. However if for any reason (such as no table space) a route cannot be stored in A, then the station will not only keep re-trying storage of a route in A, but will also keep trying storage in D. This is not ridiculously expensive since the station limits how often it is willing to retry the same route, but is clearly undesirable.

We proposed that a generalization of one of the suggestions for preventing route assignment when the use of point-to-point routes was explicitly being overridden (see E above) could solve this problem as well. A new header bit would be defined to tell whether or not a slot in the PR's route table was used in putting a route into the packet. It would be cleared by a PR which found or created a route slot for the packet route, and set otherwise. Thus it would be set for any of the following reasons: use of the route table was explicitly overridden (for example, by the PMON program in the TIU); the packet already contained a route when the PR got it from the attached device; the route in the received packet was being reversed; there was no slot for the route in the table, and none could be created because the table was full. Thus if the station sees the bit off, it means the PR referred to a slot for the route, so if the station assigns one it will use it. If the station sees the bit on, it means the PR is not using the table and there is no point in assigning a route.

2.3.6. IEEE article

Further editorial comments on the Packet Radio article for this fall's special issue of the IEEE communications were sent to ARPA; one set of comments was sent early in August, and comments on a later draft were sent at the end of August. Our concern centers on making the description of the PR net both comprehensive and reflecting the present state of the project.

2.3.7. Memory use

During this quarter we prepared a breakdown of the size of various components in each of the processes in the station. This is basically a memory utilization list, although some storage is allocated dynamically and thus a specific size for such blocks cannot be listed. The breakdown follows the explanatory paragraphs below.

Station Memory Use
August 1, 1978

(All numbers herein are decimal. All amounts of memory are in units of 16-bit words, rather than in units of 8-bit bytes.)

A PDP-11/40 can have up to 131072 words of memory physically attached to the machine. The memory management hardware, however, permits addressing at most 32768 words at any one time. Memory is further subdivided into pages of 4096 words. Thus an address space may be filled with at most 8 pages. Although the hardware permits use of partial pages, the ELF operating system does not support this. Each page must either be fully addressable in an address space, or none of it addressable.

ELF itself occupies seven pages, and the I/O registers one page, so one address space is filled with these. Other address spaces may be loaded with one or more processes. Each process must have its own stack, for which 256 words are used.

Additionally, each process may be written in BCPL, and thus require the BCPL runtime package (81 words).

One of the machine's 32 pages of physical memory is used to hold the BCPL library routines. These are shared by all processes written in BCPL and referencing the library package. The one page is entirely filled by code, one stack, and a little room for expansion. If a given process does indeed reference this library, the address space of the process will be increased to support this use. Namely, a stack (256 words) for the library will appear, as will some room for teletype buffers, variables, etc. (304 words). If, once the program begins execution, it calls library routines to initialize things for later use of the free storage package or the delayed signal facility, then further memory space will be allocated to the calling process.

Since an address space may contain fewer than 8 pages, there can be more than 4 address spaces (32 pages (system total) divided by 8). In fact, there can be several address spaces. Each address space can have several independent processes executing in it, each with its own stack, etc.; only one process in any given address space can use the BCPL library routines, however. When ELF allocates some of an address space to a process, it does so in chunks of 16 words. Thus, for instance, stacks allocated by ELF always start on 16 word boundaries. This leads to a few wasted words here and there within most processes.

- -

CONN -- 15016 total

Code	9416
Storage pool	5600

(tables and buffers are dynamically allocated from the free storage pool)

- -

GATEWAY -- 8325 total

Code	4037
Routing tables	192
Storage pool	4096

(tables and buffers, except for the routing tables, are dynamically allocated from the free storage pool)

- -

LABEL -- 24115 total, plus storage allocated at run time by/for library

Code (including parameters etc.)		20491
Buffers and IORBs		762
tell MEAS existing IDs	33	
send measurements to MEAS	62	
tell CONN routes	82	
read fwding info from CONN	17	
read/write packets	568	
Tables		2094
cumstat counters	20	
connectivity	184	
PRs	745	
non-PR devices	97	
connections	71	
measurement parameters	7	
connectivity changes for MEAS	22	
PTP routes	871	
recent PTP route attempts	77	
Process stacks		768

```
- - - - - - - - - - - - - - - - - - - - - - -
MEAS -- 24133 total, plus storage
 allocated at run time by/for library
        Code (including parameters etc.)          17522
        Buffers and IORBs                          2361
          read IDs from LABEL                 13
          iorbs for PR/TIU connections       420
          disk use                           789
          meas entry buffs and iorbs
            not included in above            969
          operator dialog                    170
        Tables                                     2008
          meas run specs                    1104
          init cumstat control pkt code       42
          connections                        155
          command tables (dialog)            707
        Text strings (other than in above command tables)
                                                   1218
        Process stacks                             1024
- - - - - - - - - - - - - - - - - - - - - - -
STACON -- 6362 total
        Code                                  3347
          STACON                      2907
          TELNET                       440
        Storage                               3015
          STACON                      2739
          TELNET                        20
          stack                        256
- - - - - - - - - - - - - - - - - - - - - - -
TCP -- 4283 total
        (memory use breakdown not available at this time)
```

- -

XRAY -- 3889 total

Code		2147
Tables		436
parameters, etc.	79	
canned text	357	
Variables		409
packet buffers	280	
other variables	129	
Environment overhead		897
BCPL runtime code	81	
process stack	256	
library stack	256	
library variables	304	

- -

2.3.8. Robustness meeting

During this quarter there has been rising concern over the performance of the PR net, and specifically over the reliability and robustness of the network components. This concern culminated in a meeting at SRI in August. In two communiques we suggested a number of topics which could profitably be addressed at this meeting. These topics are:

A. An experiment should be designed, executed and analyzed to assess the true impact of tying up one of a PR's buffers. This follows up data reported by SRI during the PRWG meeting at BBN.

B. A simple experiment should be made to obtain data on how often, and for how long, the station PR blocks traffic for it from the station. This would allow assessment of whether this potential difficulty is contributing significantly to the problems observed in practice (such as loss of CUMSTATs).

C. Further study should be performed to evaluate whether modification to connection process logic, to permit sending of acknowledgments for duplicates of previously accepted packets, would help significantly.

D. Besides study of station memory limitations, planned for later this year, we suggest systematic experiments to search for values of the retransmit count and interval for SPP in PRs, offering better performance.

E. Additional consistency checks on packets arriving at the station, beyond those already in existence, can be implemented; network operations policy can be tightened up to investigate anomalous conditions when detected by such checks.

F. Follow up existing evidence in the missing CUMSTATs problem.

G. PR diagnostics, and the use of them, should be examined to identify areas needing improvement. The diagnostics should be updated as new failures appear which diagnostics could have recognized but didn't. The Load-Verify PR operating system command could be changed to continue comparisons and printout of discrepancies beyond the first mismatch; this would aid debugging and failure analysis.

H. Provision for making various consistency checks on the network, both on-line and off-line, could be enhanced considerably. More data could be captured when a malfunction becomes apparent. The usefulness of such measures should be considered, as well as their costs. BBN suggested seven candidate measures of this sort.

I. The connectivity "gap" problem should be pursued.

We told SRI how to modify the parameter in the connection process which controls the maximum number of retransmissions of an SPP packet. They needed more retransmissions than normal for some UCLA experiments in which the station was failing to get ACKs and thus aborting measurement connections.

During this quarter Collins released PRTN 257 on the down-line loading of present PRs (EPRs). We discussed their

design and expect to forward to them a number of comments early in the next quarter. Our concern is that some aspects of the down-line loading mechanism are not specified fully, so it is unclear just how it will work as a whole.

3. THE PACKET RADIO NETWORK

3.1. Station Programming and Testing

3.1.1. ELF system

A new library was installed on both SRI machines this quarter, which fixes the problem with the ELF clock jumping ahead when the hardware clock count wraps around. This concludes work on this problem begun in preceding quarters. We also designed a change to the station debug process to permit one experimenter to take over the debugging of a process previously being debugged by someone else. This addresses a deficiency found in the visit to SRI for point-to-point tests this quarter, in which a debugging connection failed and could not be reestablished due to the previous exclusiveness check.

3.1.2. Labeler

A major accomplishment this quarter was the release of station software which supports point-to-point routing. In this section we describe the software changes as of this release and discuss some of the station testing which was carried out.

The only station programs that were modified were the connection process and the labeler. The changes included bug fixes and performance improvements as well as the new point-to-point routing. A summary of changes follows.

Connection process

1. The connection process includes the forwarder (for intranet packets). An interface between the forwarder and the labeler was implemented, whereby the forwarder tells the labeler the source and destination of packets it is forwarding. The labeler thus knows what point-to-point routes are needed.

2. Unneeded buffers are released sooner than before when closing connections. This should make it much less likely that the

measurement process will fail to close connections to PRs or TIUs from which it is collecting measurements. The failures, which were quite common, resulted from a lack of connection process buffers to use to generate the FIN packet for the connection.

3. A bug was fixed that allowed transmit windows to become zero. This prevented the nominal retransmit interval from being used, and in fact retransmissions were significantly delayed. Hence, for example, aborting a connection due to lack of an ACK took too long. This bug explained many mysterious labeling delays we had noticed in measurement files, which had been troubling us for some time.

4. A change was made in the generation of ACK sequence numbers. Formerly ACKs for duplicates all used the same sequence number and could get erroneously filtered in PRs as if they were network duplicates.

5. If no route is known to put in an ACK, the ACK is sent by reversing the packet route. Formerly if no route was known no ACK was sent.

6. If no buffer is available to ACK a packet, the packet is accepted and the failure to ACK is ignored. Formerly a warning message was printed and the connection was aborted.

Label process

1. Point-to-point routes are assigned. Commands are provided to enable the station operator to display current point-to-point routes and to control whether or not point-to-point routes are given out. The labeler's handling of point-to-point routing is described in PRTN 174, revision 6.

2. Routes which pass through terminal and station PRs are assigned if necessary. Formerly routes were assigned only through repeaters.

3. In order to help identify network problems, such as PR hardware problems, labeler warning messages (messages which are printed when some packet fails a consistency check) were expanded to include the entire packet contents. Also more consistency checks were added to ROP processing.

4. Several labeler parameters were made settable by operator command. Formerly these could only be modified using the cross-net debugger. This should improve SRI's ability to experiment with network behavior. Two key parameters which can be set are the connectivity refresh interval, which determines how long the labeler awaits a ROP before erasing connectivity, and the time after which a point-to-point route will be considered old and be erased.

5. Various error conditions cause messages to be printed on an operator console. An option was added to make the labeler halt on errors indicating station bugs, to improve the chances of finding the bug.

6. A slight change to ROP interpretation was made in preparation for the CAP4.8 change in DROP broadcasting. This prevents the DROPs from being misinterpreted as implying connectivity from a PR to itself.

Before delivering the station, we tested it extensively in the SRI net. Testing was spread over a long period due to unavailability of SRI facilities caused mostly by PR hardware work and demos. Most of the testing was done remotely, from BBN. We tested the basic functioning of labeling and point-to-point route assignment (between two TIUs in the lab with direct connectivity) during SRI's off hours, with no SRI personnel in attendance. We did mobile testing while SRI personnel were en route to mountain repeater sites and the airport in the packet radio van, since our running the new software did not interfere

with what they were doing. This testing exercised assignment of routes through terminal PRs and assignment and reassignment of point-to-point routes between mobile TIUs in the van and a TIU at SRI. Two problems of interest, which still need to be dealt with, showed up during these tests. One was the interference of point-to-point route assignment with the TIU's PMON program. This program sends packets to itself to measure round-trip delay through the station forwarder, but with point-to-point routes, the packets simply looped back in the TIU's own PR. This problem and solutions are discussed fully in section 2.3 above. The other problem was a subtle station bug, which we were not able to track down. This did not preclude delivery of the new software, however, since it had shown up at least once before, with the old software. We still hope to debug it, and have instructed SRI on how best to save the evidence for us if it happens again.

In addition, we traveled to SRI for a few days of on-site network testing. Although we did use the new station software during these tests, and did find a bug in it, there was nothing involved in testing the point-to-point software that could not be carried out identically from BBN. The primary purpose of the visit was to investigate the loss of PR measurement data on SPP connections (see section 3.2), labeling delays (see connection process item 3 above), and gaps in response that SRI had observed several times during van rides. For all of these situations it seemed that the packet records collected by SRI's Interdata packet monitor might help shed light on what was happening. Unfortunately, the visit was plagued by failures of PRs and service host machines, drastically limiting how much could be accomplished. In one case, the Interdata PR failed, so that packets at the data rate of interest were not recorded. Another run produced no data because the mobile PR failed. Unavailability of ARPANET service hosts made it impossible to

collect, format, and print measurement data from the runs that were successfully carried out until we had returned to BBN. Printouts of the Interdata packet records for the times of interest were not available until a month after the testing, because the printer at SRI was down. We have recommended that steps be taken to make the Interdata records, currently recorded on a magnetic tape, available on-line, on a service host. Currently all that can be done with the tape is to print information in hexadecimal about all packets heard. This makes it almost impossible to study more than a minute or so of data. With on-line availability, programs could be written to reduce the data and extract useful information in a usable form.

3.2. Network Performance Investigation

At the June Packet Radio meeting, UCLA reported preliminary data from their first real experiment at SRI. The experiment called for CUMSTATs (cumulative statistics) from several PRs and several TIUs. Each PR and TIU is supposed to send CUMSTATs to the station at regular intervals, but UCLA's data included many fewer CUMSTAT packets than expected. CUMSTATs are delivered to the measurement process in the station using the SPP protocol to provide reliable delivery. The packets are retransmitted a number of times, until an end-to-end acknowledgement is received; if none is received, the connection is aborted and the PR or TIU sends no further CUMSTATs. There had been 51 experiment runs, each lasting 11 minutes and requiring a CUMSTAT once every 2 minutes from each of 4 PRs. 32 runs out of the 51 failed to collect the expected number of CUMSTATs on at least one connection. 58 out of the 204 connections delivered fewer than the 5 expected CUMSTATs. CUMSTATs were also collected from TIUs, but UCLA did not look at whether there were problems with them too. This high failure rate certainly indicated a serious

problem. Since CUMSTATs are collected by the station, which we programmed, and are delivered using SPP, which we designed, we were called upon to investigate this problem. The possible causes of the problem seemed to fall into three categories: (1) SPP specification (inadequacies in the protocol design), (2) implementation errors (incorrect programming in the PR, TIU, or station), and (3) network performance (failure to deliver packets after the allowed number of retransmissions).

Over the course of several weeks we requested and received from UCLA detailed data on the failed connections. This took a while because we did not know what further questions to ask until we had studied what they had already sent us, and they had to write programs to extract the desired information from the measurement files. The following information was collected from UCLA and analyzed:

- Number of packets received from each PR in each run
- Whether CUMSTAT packets in the file were retransmissions (they never were!)
- Whether timestamps indicated correct intervals between packets in the file
- Breakdown of packets received from each PR in each run by SPP function (open connection, normal packet, close connection, open-and-close)
- What happened from the measurement process's point of view when it tried to close the connections at the end of the run

Based on this information, we concluded that there was probably a malfunction in the PRs. That is, there were symptoms that could not be explained if the PRs were handling connections correctly. Further evidence for PR malfunction showed up during some station testing we did at SRI. We were collecting CUMSTATs from the station PR at one-minute intervals, but found that they

only came in for the first 25 minutes (out of 50 minutes total). The station PR should certainly not fail to deliver its packets to the attached station, and the symptoms when the station tried to close the connection were inconsistent with the PR having given up due to lack of acknowledgement. We worked with Collins to help them debug the PR code. Two bugs were found: One caused a PR to generate a CUMSTAT with a garbage route, which would not be delivered, and the other caused the PR to not abort the connection properly upon CUMSTAT delivery failure. A repeat of UCLA's experiments with the corrected PR code exhibited almost no missing CUMSTATs.

3.3. Support

Several changes were made to the PRDATA program, which is used to print measurement data collected in the Packet Radio Network in a form suitable for human processing. The changes were to print point-to-point route entries made by the new labeler, to fix several bugs, and to aid in analysis of labeling activity.

The labeling analysis changes were developed during the labeler testing reported in our last quarterly report, in which we studied the effects of various parameters on the amount of relabeling. They automated some of the common things we were originally doing by time-consuming visual scanning, thus greatly increasing the amount of information available to us. The measurement file contains entries made by the labeler whenever it relabels a PR or fails in an attempted relabeling. In order to help us see what a labeling change was about, we changed PRDATA to keep track of the latest labeling as it processes the file, and mark each labeling entry it prints to show how this labeling (route and good neighbors) differs from the previous values. PRDATA also prints references to previous labeling entries which

indicated failed attempts, so as to show how the current entry confirms success or failure of the previous attempt. At the end of the measurement report, PRDATA prints a summary of labeling activity seen in the file, showing for each PR how many successful or unsuccessful labeling attempts there were, and what changes (route and/or neighbors) were made during relabeling.

Here are some extracts from a PRDATA printout, showing these new features. The numbers in the left margin indicate elapsed time in seconds since the start of the measurement run. Numbers in brackets are the contents of the entry in hexadecimal, and can be ignored for our purpose here. Asterisks are used to indicate new good neighbors. The numbers in the summary may fail to add up because the summary does not count labeling of a PR which is unlabeled or reports bad labeling as a labeling change.

```
 3.074   LABEL: PR Labeling [10]
     [900]      PR ID = 8016   Valid: PR labeled as follows
     [903]      Route word
     [0]  Route word
     [0]  Route word
          Route to station: A020<--8016

     [A04]      Neighbors: 8014, 8023
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)

34.745   LABEL: PR Labeling [10]
     [900]      PR ID = 8016   Valid: PR labeled as follows
     [903]      Route word
     [0]  Route word
     [0]  Route word
          Route to station: A020<--8016

     [A03]      Neighbors: A020*, 8023
     [C]  Neighbor:  8027*
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
          (Erased:8014)
```

```
125.364   LABEL: PR Labeling [10]
     [900]     PR ID = 8016   Valid: PR labeled as follows
     [A03]     Route word
     [9]  Route word
     [0]  Route word
          Route to station: A020<--8023<--8016
          (Changed from:    A020<--8016)

     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
          (Erased:A020,8023,8027)

485.223   LABEL: PR Labeling [10]
     [503]     PR ID = 9012   Valid: following label not ACKed
     [403]     Route word
     [506]     Route word
     [0]  Route word
          Route to station: A020<--8014<--8013<--9012
          (Changed from:    A020<--8027<--9012)

     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)

495.617   LABEL: PR Labeling [10]
     [500]     PR ID = 9012   Valid: PR labeled as follows
     [C03]     Route word
     [5]  Route word
     [0]  Route word
          Route to station: A020<--8027<--9012

     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     [0]  (Empty Neighbor Slot)
     UnACKed label report at 485.223 failed
```

LABELING SUMMARY

PR	succeeded with ACK*	failed	succeeded but not ACKed	changed route & neighbors	changed route only	changed neighbors only
9011	14	0	0	0	13	0
9022	1	0	0	0	0	0
A020	1	0	0	0	0	0
8014	12	0	0	4	0	7
9012	11	1	0	0	9	0
8013	11	0	0	2	0	8
9005	1	0	0	0	0	0
9021	1	0	0	0	0	0
8016	6	0	0	3	0	2
8023	10	0	0	2	0	7
8015	1	0	0	0	0	0
8027	9	0	0	2	0	6
	78	1	0	13	22	30

*(success count biased by inclusion of initial labeling reports)

4. INTERNETWORKING

Research, development, testing and support activities in the field of Internet communications are a significant part of our efforts under the Packet Radio contract. During this quarter, particular progress was made in the areas of Transmission Control Protocol, gateway routing, and Very Distant Host interface performance. Each of these is discussed in detail in the subsections below. Our work on the Host/SIMP protocol module for ARPANET/SATNET gateways and efforts on the minigateway, both using significant manpower resources in previous quarters, have been on a back burner this quarter, so to speak. The final debugging of Host/SIMP awaits availability of an actual SIMP loaded with software implementing the protocol already agreed upon and supported by the existing code in gateways. This SIMP software is being implemented by another BBN group under separate contract to ARPA; delivery is anticipated next quarter. The gateway's Host/SIMP module has been tested as fully as practical without a true SIMP, as described in the previous quarterly report. The minigateway software, working well as discussed in last quarter's report, awaits actual LSI-11 minigateway hardware from DEC, SRI and ACC before final testing and delivery is possible.

4.1. Transmission Control Program (TCP)

During June, 1978 a second release of the hand coded TCP (for the TENEX and TOPS20 monitor) was given to SRI for installation on the SRI-KA host. The main differences from the first release were:

- o Code for supporting obsolete protocol features (INT, RSN, ARQ, etc.) was removed.

- A bug which caused data to be dropped when multiple connections were in use was fixed.
- The TENEX 96-bit leader IMP driver module was used.

A compatible version of the BCPL TCP was constructed and installed at SRI-KA at the same time. Also, TTLSRV (the TCP TELNET server program) was modified so that it could be run under either the new monitor TCP or the BCPL TCP without reassembly. This version of TTLSRV along with changes made to the TIU by SRI supports eight-bit streams and permits various graphics editors to be used over TCP TELNET connections.

After installation of the TCP but prior to the July 11, 1978 BAA demonstration, two more bugs were found and repaired. One of these had to do with code which attempted to minimize the amount of free storage tied up to hold packets awaiting reassembly which had become subsets of "bigger" packets (retransmissions from a TIU). The second major bug had to do with the address matching routine which was finding old, dead connections before they had been garbage collected. Before fixing this, it was sometimes impossible to reopen a connection without waiting for thirty seconds.

Somewhat later a bug was fixed which manifested itself by turning off the retransmissions from the TCP. Also, SRI provided a monitor crash file. The TCP part seemed to be intact but various other parts of the monitor were full of garbage. The cause could not be identified but a hardware malfunction was suspected.

The TCP source files have been modified so that they may be assembled as part of a TOPS20 Release 3 monitor which will run on a 2020 machine. Some time has become available on a 2020 at BBN (until August 11) and it is hoped that this time can be used for active TCP debugging.

4.2. Gateways - Routing and VDH Work

We began coding and debugging of the alternate routing strategy in gateways as outlined in IEN #30, "Gateway Routing, An Implementation Specification". We also coded and debugged a version of the gateways that supports 96-bit leaders on the ARPANET. We made some modifications to the ARPANET/SATNET gateways, primarily to make these gateways more robust. The gateway now checks message lengths on network interfaces to insure that a network does not attempt to send or receive larger packets than the specified maximum for that network.

A considerable amount of time was spent this quarter in attempting to understand and improve the gateway's performance in handling the VDH interface. Tests performed with the SATNET SIMPs indicated that many messages were being retransmitted over the VDH interface, thus contributing significantly to the delay seen by messages in traversing the SATNET. This problem was traced to the method in which gateways handled input from the VDH. The VDH interface has two channels to provide buffering on the interface. Although these channels must be used in turn, i.e., the first packet must be sent on channel 0, the second on channel 1, the third on channel 0, etc., a VDH interface should be prepared to accept packets on either channel and buffer them for later reassembly into messages for the user. The gateway interface was refusing traffic received out of order, thus accounting for the large number of retransmissions and associated delay on the interface. This problem was corrected and we verified that retransmissions on the interface were substantially reduced. Despite these modifications to the VDH interface, throughput was not increased on this interface. Our tests showed that the ELF operating system responded slowly to interrupts and was very nearly processor bound when sending traffic at the rate

of 8 packets per second. Thus, the poor throughput on the VDH interface was not due to any errors in the modules handling the interface, but rather to the limitations of the operating system. These results were reported to the SATNET Working Group and the modified versions of the gateway code were installed in the ARPANET/SATNET gateways.

5. HARDWARE

5.1. Boston Area PRN Site Preparations

During the month of June we began to prepare the facilities at BBN for receiving and installing the two PR units due to arrive in July. This included providing the room each unit will be located in with whatever is necessary for operation of the unit; determining feasible paths for running the 1822 interface and the coaxial cables; choosing and ordering the necessary hardware (i.e., cables, connectors, cabinets); and constructing special equipment. These are described in detail below.

5.1.1. PR unit locations

A small room on the top floor, earmarked for Packet Radio use since the construction of BBN's facility at 10 Moulton Street, was given over to us as of July 1st. It will house the analog and digital components of one unit, plus a terminal. (The antenna will be outside.) Electricity, necessary for turning on the power supply, was lacking so two outlets were installed. A metal junction box was also installed to provide for splicing the distant host burial cable to more flexible cable before plugging it into the PR unit.

Consideration also had to be given to whether additional air-conditioning should be provided in the room. It was estimated that the electrical equipment would dissipate approximately 100 watts, therefore no additional air-conditioning was needed.

The room on the first floor (North Bay), where the second unit is to be located, needed no further preparation. Everything necessary for installation was present, and there appeared to be no conditions that might interfere with successful operation of the unit.

5.1.2. Cable paths

Paths for two cables had to be determined: one for a coaxial cable to connect the Packet Radio unit on the seventh floor to the antenna on the elevator shaft, and a second for a direct burial cable to connect the unit to the host computer on the ground floor. It was also necessary to use no more than the maximum cable length recommended for the coax cable, else the decibel loss would be too large for satisfactory RF transmission and reception. Maximum length for the coax cable, SRI informed us, is 60 feet.

The coax cable, it was found, could go through the seventh floor's ceiling and along the false ceiling to a pipe, which had been installed expressly to provide passage for the cable through the concrete roof. Emerging from the pipe onto the roof, the cable could then go over to the nearby elevator shaft, then up the shaft to the antenna.

The direct burial cable, it was found, could go along the false ceiling above the first floor to right underneath the second floor's telephone cable room. From there, it could go up to the sixth floor through pipes in these phone rooms; these rooms lie directly on top of one another, and each has a large pipe through its floor to receive the phone wires from the floor below. At the sixth floor, the cable could exit the phone room via an existing hole in the wall. At this point developed the only real problem we had with this cable path. It is only the first through sixth floors which have connected phone rooms; the other floors are solid concrete. Once on the seventh floor, the cable could go up into its false ceiling and along to the Packet Radio room, then down into the room. But there seemed to be no prearranged pipes to take the cable from the sixth floor to the seventh. Happily, a drainpipe hole was found. Had nothing been

found, it would have been necessary to drill a hole, an expensive and potentially hazardous operation.

5.1.3. Hardware ordered

We were not sure what type of cables to order or from whom to order them. At our request, SRI sent us the name, manufacturer, and part numbers for suitable coaxial and distant host cables and connectors, and we ordered these. It also seemed to be a good idea to plan on splicing the host cable to a more flexible cable, and plug that into the PR unit, thereby lessening the chance of the cable ripping out of the unit if either unit or cable was moved or stressed. Cable for this purpose was ordered as well.

This equipment was ordered early in July. The host cable was scheduled for delivery on August 7, and the other cables and connectors, on September 1. The latter items arrived during the last week of August. However, the host cable arrived several weeks later than scheduled, and had not arrived by August 31, the end of the time period covered by this report.

Other hardware included enclosures for the two units. We decided to purchase from Digital Equipment Corporation one short (height 50 inches) cabinet with casters, and to assemble a second cabinet from surplus parts that Packet Radio's division already owns.

The short cabinet would enclose the PRU on the first floor. We opted for a short cabinet for this unit so that a terminal could be placed on top. This combination of short cabinet and casters offers three advantages: the unit can be moved around the North Bay during testing; it can be moved easily because of its compactness, with both the PRU and terminal mounted on one piece of equipment; and the operator is likely to walk back and

forth between the unit and the host, so stand-up typing would be convenient. Additionally, these features will make this unit an attractive one to show visitors.

The second cabinet, which would enclose the PRU on the top floor, will be of standard height. It will also be on casters, though this is simply for convenience; the unit is not intended to be a mobile one. The terminal will be placed on an adjacent table.

The short cabinet was ordered in mid-July, arrived about five weeks later, and has been assembled. Parts for the second cabinet have also been gathered and assembled. The units will be mounted in the near future.

5.1.4. Construction of special equipment

Since we wished to separate one antenna from its unit for mounting outside on a brick elevator shaft, we asked SRI whether the brackets provided with the unit would be suitable or whether something else was needed. SRI responded that the brackets were not adequate and a special apparatus was needed. Together with Collins they drafted plans for such an apparatus, and sent the plans to us. The plans called for construction of a mast, or hollow tube; a collar, or flat donut ring; and two mounting, or right angle, brackets, all to be made of aluminum. When in place, the antenna is bolted to the collar, the collar screwed onto the top of the mast, the mast clamped to the brackets with ring clamps, and the brackets are bolted to the elevator shaft.

These parts were built and secured to the elevator shaft by the end of July. By the end of August the antenna itself was mounted, and the coax cable threaded through the mast and connected to the antenna.

5.1.5. Problems encountered

Preparations proceeded smoothly for the most part. A few problems were encountered, but none constituted a serious setback. The first was a delay in delivery of the 1822 interface cable. Delivery was scheduled for the first week in August, but the cable had not yet arrived by the end of August. The manufacturer had named a date for expected delivery to the vendor, but found itself unable to complete the order by that time.

A second difficulty arose concerning the antenna-mounting apparatus. The various components were built according to the specifications given us. However, when the attempt was made to bolt the antenna to the collar, it was found that the four holes on the collar did not align with the holes at the base of the antenna. Re-examination of the collar construction plans showed that no measurement had been given to indicate how far in from the edge of the collar to place the tapped holes, and so this distance had been guessed at. To correct this, the proper measurement was recorded and, after turning the collar a few degrees on its axis, new holes were tapped at the correct distance.

A third difficulty concerned installation of the coax cable. This cable was stiffer than we had anticipated, so some adjustments had to be made to permit the cable to be laid along its path to the antenna. The pipe that had been installed to provide passage for the coax cable through the ceiling to the roof had a curved head. However, this head was curved too sharply to allow the stiff cable to pass through. To remedy this, the head was removed. It was not replaced with a head of a different shape, though. Instead, weatherproof sealing paste was put thickly around the edges of the pipe and cable. The absence

of a curved head eliminated the need for bending the cable, and it is expected that the paste will provide adequate weather protection.

The antenna mast also required an adjustment because of the cable's rigidity. The mast is attached to the elevator shaft close to where the pipe goes through the roof. The bottom of the mast was too close to the pipe to allow the cable to exit from the pipe and then enter the mast. Therefore, several inches were cut off the bottom end of the mast to allow the cable to approach the entrance more gradually.

5.1.6. Additional PRs

Preparations have begun for next year's installation of Packet Radio units at M.I.T. and at Lincoln Labs, and of a repeater on Zion Hill. In response to inquiries about our own preparations, copies of our purchase orders for cables and connectors were sent to both Lincoln and MIT, giving the manufacturer and part numbers. We will keep them informed of our progress, and will try to answer questions they may have as their own preparations progress.

We have also obtained some pricing information for leasing space on one of the Zion Hill radio towers. However, formal negotiations for this have not yet begun.

5.2. Miscellaneous Hardware Work

The information on mini-gateway hardware configuration, in preparation at the close of last quarter, was sent to SRI. In return, we received a description of their hardware plans for Terminal Interface Units (TIUs) and port expanders (for host ports on ARPANET IMPs). Their plans mesh well with ours, and we have decided to rely on their hardware in two major respects. First, the SRI/ACC LSI-11 1822 interface will be used instead of

the Collins 1822 interface. Second, the SRI Robustness module will be used instead of the MRV11 memory board with ultraviolet EPROM and the special restart-on-HALT card we had previously planned to build at BBN. These design modifications will, we believe, save money, provide interchangeability with TIUs and port expanders, and secure somewhat superior operation than possible with the earlier configuration. We also communicated cost information on the hardware components to ARPA where possible (the parts from SRI and ACC have costs to be determined between ARPA and those suppliers).

During this quarter a thin film RAM memory board failed in one Packet Radio Digital Unit. The failure resulted in destroyed code in a few, seldom-used locations. The failed board was returned to Collins, and they sent a working replacement.